Watters, A. (2014). The hidden history of ed-tech. In The monsters of education technology (pp. 7-31). Vancouver, BC: Tony Bates Associates Ltd. Retrieved from https://s3.amazonaws.com/audreywatters/the-monsters-of-education-technology.pdf#page=7

Watters, A. (2014). The hidden history of ed-tech. In The monsters of education technology (pp. 7-31). Vancouver, BC: Tony Bates Associates Ltd. Retrieved from https://s3.amazonaws.com/audreywatters/the-monsters-of-education-technology.pdf#page=7

EDDL 5101
Educational Technology for Learning

I. THE HIDDEN HISTORY OF ED-TECH

1 THE HISTORY OF THE FUTURE OF ED-TECH

A couple of months ago, my brother Fred and I went back to the house we grew up in. We're getting to the age where we have parents that are "that age." Our dad had fallen, broken his hip, and was in a nursing home. We went to his house to "check on things."

It's been over 20 years since either of us lived there. My bedroom has since become the guest bedroom. With the exception of the old bookshelf and bed, there's nothing there that's "me." But my brother's room has become a shrine to the Fred that was. It remains almost entirely untouched since he graduated from high school and moved out to attend the Air Force Academy. It's pretty weird to visit the room now. Fred didn't stay in the Air Force Academy. He dropped out after his sophomore year, became an environmental activist and then an emergency room nurse; he's now a nurse practitioner in Maine.

Visiting his old bedroom was like stepping into the past that felt strangely disconnected from the present. Not totally disconnected; strangely disconnected. You could find glimpses there of the kid he was, of the person he was supposed to become – of my parents' and grandparents' visions of and plans for his future.

A future predicted in the 1970s and 1980s.

It's not the future that came to be. The room contained the history of that not-future.

We found on his bookshelf another example of this: *The Kids' Whole Future Catalog*. Published in 1982, I remember how Fred and I would pour over this book. My brother admitted that it had shaped his thoughts at the time on who he would become, his expectations of what his future would look like.

It's a future of food factories and space vacations and flying trains. It's a future where robots are teachers.

Again, like my little brother's room, it's mostly the future that wasn't. Or the future that isn't quite. Or the future that isn't quite yet.

I want to talk about the history of the future of education technology this morning. The future that wasn't. The future that isn't quite. The future that isn't quite yet. The history of all of this.

I want us to consider where education technology has come from, where we have been. As educators. As technologists. Where are we going? What narratives help us answer questions about the past? What narratives help us shape the future?

As we move forward into a world that it increasingly governed by machines and algorithms, I think we must consider the trajectory of the path we're on. Because "the future of ed-tech" is shaped by the history of ed-tech – whether we realize it or not.

Last year, the programmer and designer Bret Victor delivered what I thought was one of the smartest keynotes I've ever seen.

He came on stage dressed in a button-up shirt and tie with pocket protector and proceeded to utilize an overhead projector and transparencies for his talk. There are three visual cues there about the conceit of his talk: the pocket protector, the overhead

projector, and transparencies.

Victor spoke about the future of programming, but as those visual cues and those presentation technologies implied, he spoke of the future as though he was describing it in 1973. "Given what we know now," he asked, "what might programming be like 40 years from now?" In other words, given what we knew about computing in 1973, what would it be like in 2013?

Victor proceeded to talk about some of the computer science research that had been conducted in the previous decade – that is, in the 1960s. This was the research that he used to inform his predictions about the future:

Gordon Moore, Intel's co-founder for example, postulated in 1965 what we now call "Moore's Law," the observation that over the history of computing hardware, the number of transistors on integrated circuits doubles approximately every two years. In other words, the processing power of computer chips doubles roughly every two years. This is a prediction that has come true, but in part because the technology sector was worked to make this a self-fulfilling prophesy of sorts. Chip manufacturers like Intel have made increased computing power year-over-year an explicit goal.

But much of the research that Bret Victor cites in his keynote was never really adopted by the technology industry. It simply wasn't the focus. It wasn't the goal. Victor points to a number of incredibly interesting and provocative research efforts that went nowhere.

The Future of Programming that wasn't.

And it isn't that these innovative ideas were rejected or ignored because they just weren't do-able. What's worse, they were do-able, but they were ignored and forgotten. The technology that powers our computing systems today took a very different path than the one that Victor wryly describes in his talk. And today, many programmers don't recognize, let alone teach others, that there could be other ways of doing things, of designing and developing technologies.

Take the work of Douglas Englebart, for example. He passed away last year, an amazing but largely unsung visionary in computer science. Among other things, Englebart was the first to use an external device that rolled around on a flat surface and moved a pointer on a screen so as to highlight text and to select options – what we now call the mouse.

Englebart unveiled the mouse in what technologists refer fondly to as "The Mother of All Demos," a demonstration in 1968 of the oN-Line System (more commonly known as NLS), a hardware and software package that had a number of incredible features demonstrated publicly for the first time. Again, remember, this was the era of the mainframe and the punch-card. In the demo: the mouse, "windows," hypertext, graphics, version control, word processing, video conferencing, and a collaborative real-time editor.

1968.

But many of the features in "the Mother of All Demos" weren't picked up by the tech industry – at least, not right away. The team that had worked with Englebart on the NLS soon dispersed from their Stanford University-based research program, many of them ending up at Xerox PARC (Xerox's Palo Alto Research Center). In turn Xerox PARC became the site where many more of the computing technologies we do now take for granted were developed, including the Ethernet, laser printing, and the personal

computer as we know it today.

But even at Xerox PARC, new technologies were devised that were never widely adopted. Why? Why, when as Victor argues, many of these were more interesting and elegant solutions than what we have actually ended up with?

In part, it's because computing technologies can be prototyped quite readily. The 1960s and 1970s marked the beginning of this: computers had become powerful enough to do interesting things, and many computer scientists were busily imagining what those interesting things might be and how they might be done. But while building new technologies is easy (or easy-ish), changing behaviors and culture is much, much harder.

What does this have to do with ed-tech?

Well, the tension between new tools and old practices should give you a hint. It's simple to introduce iPads into the classroom, for example. It's much more difficult to use them to do entirely new things, particularly things that run counter to how classrooms have operated in the past.

Watching Victor's talk, I couldn't help but wonder how might we have written "The Future of Ed-Tech" in the 1970s. After all, Bret Victor says that his keynote was inspired by Alan Kay, an important figure not just in programming but in education technology as well.

Kay was a pioneer in object-oriented programming. He actually attended Englebart's demo in 1968, and he later worked at Xerox PARC where he helped develop the programming language SmallTalk. (MIT Media Lab's introductory programming language for kids, Scratch, is based in part on SmallTalk.) And Alan Kay designed the prototype for something called the DynaBook, "a personal computer for children of all ages."

If I were to tell you the story, using the conceit that Bret Victor used in his keynote – that is, if I were to come out here today and tell you about the future of education technology as it might have been seen in the early 1970s – I would ground the talk in Alan Kay's DynaBook.

Again, let's recall that in the late Sixties and early Seventies, computers were still mostly giant mainframes, and even the growing market for microcomputers was largely restricted to scientists and the military. Alan Kay was among those instrumental in pushing forward a vision of personal computing.

We scoff now at the IBM CEO who purportedly said, "I think there is a world market for maybe five computers." But "personal computing" for Kay wasn't simply that computers would be adopted in the workplace. That's something you can imagine that every IBM executive would readily agree to.

Kay argued that computers should be commonplace and be used by millions of non-professional users. Kay believed this would foster a new literacy, one that would change the world much like the printing press did in the 16th and 17th century. And key: children would be the primary actors in this transformation.

In 1972 Kay published a manifesto, "A Personal Computer for Children of All Ages," in which he describes the DynaBook, the underlying vision as well as its technical specifications: no larger than a notebook, weighing less than four pounds, connected to a network, and all for a price tag of $500, which Kay explains at length is "not totally outrageous." ($500 was roughly the cost at the time of a color TV.)

"What then is a personal computer?" Kay writes. "One would hope that it would be both a medium for containing and expressing arbitrary symbolic notations, and also a collection of useful tools for manipulating these structures, with ways to add new tools to the repertoire." That is, it is a computer program but one that is completely programmable.

"It is now within the reach of current technology to give all the Beths and their dads a 'DynaBook' to use anytime, anywhere as they may wish," Kay writes in his 1972 manifesto. 1972 – 40 years before the iPad. "Although it can be used to communicate with others through the 'knowledge utilities' of the future such as a school 'library' (or business information system), we think that a large fraction of its use will involve reflexive communication of the owner with himself through this personal medium, much as paper and notebooks are currently used." The personal computer isn't "personal" because it's small and portable and yours to own. It's "personal" because you pour yourself into it – your thoughts, your programming.

So, if I were to tell you a story about the future of ed-tech like Bret Victor tells about the future of programming, I'd probably start from there, from the DynaBook's design in 1972. And it would be a story, like Victor's, with a subtext of sadness and loss that this is not what history has given us at all.

In some ways, the DynaBook does look a lot like our modern-day tablet computers. It looks a lot like the iPad even. (Kay did work at Apple, I should note, in the 1980s under then CEO John Scully). But as Kay has said in recent interviews, the iPad is not the actualization of the DynaBook.

He told *TIME* magazine last year that the primary purpose of the DynaBook was "to simulate all existing media in an editable/authorable form in a highly portable networked (including wireless) form. The main point was for it to be able to qualitatively extend the notions of 'reading, writing, sharing, publishing, etc. of ideas' literacy to include the 'computer reading, writing, sharing, publishing of ideas' that is the computer's special province. For all media, the original intent was 'symmetric authoring and consuming'."

Consumption and creation: that's a tension that's plagued the iPads since they were unveiled. But it's one that the DynaBook was designed to balance.

"Isn't it crystal clear," Kay continued in his *TIME* interview, "that this last and most important service [authoring and consuming] is quite lacking in today's computing for the general public? Apple with the iPad and iPhone goes even further and does not allow children to download an Etoy made by another child somewhere in the world. This could not be farther from the original intentions of the entire ARPA-IPTO/PARC community in the '60s and '70s."

For Kay, the DynaBook was meant to help build capacity so that children (and adults too) would create their own interactive learning tools. The DynaBook was not simply about a new piece of hardware or new software, but about a new literacy, a new way of teaching and learning. And that remains largely unrealized.

Again, as Bret Victor's talk reminds us: changing technology is easy; changing practices, not so much.

Alan Kay's work draws heavily on that of Seymour Papert. (Bret Victor's work does too, I should add. As does mine.) Kay cites one of Papert's best-known lines in his manifesto: "should the computer program the kid or should the kid program the

computer?"

Kay's work and Papert's work insist on the latter.

Kay met Papert in 1968 and learned then of Papert's work on the Logo programming language.

As a programming language, Logo not only helped teach children programming concepts but also helped develop their "body-syntonic reasoning." That is, Logo – and particularly the Turtle that the language became most synonymous with – helped give students an embodied understanding of mathematics. There was a Turtle robot and later a Turtle graphic on the screen. Using Logo, students could manipulate these; and this, Papert argued, would help them to understand and reason mathematically.

Computers, argued Papert, should unlock children's "powerful ideas." That's the subtitle to his 1980 book *Mindstorms*, a book that both Bret Victor and I insist you read. The book addresses, "how computers can be carriers of powerful ideas and of the seeds of cultural change, how they can help people form new relationships with knowledge that cut across the traditional lines separating humanities from sciences and knowledge of the self from both of these. It is about using computers to challenge current beliefs about who can understand what and at what age. It is about using computers to question standard assumptions in developmental psychology and in the psychology of aptitudes and attitudes. It is about whether personal computers and the cultures in which they are used will continue to be the creatures of 'engineers' alone or whether we can construct intellectual environments in which people who today think of themselves as 'humanists' will feel part of, not alienated from, the process of constructing computational cultures."

Computers, Papert insisted, will help children gain "a sense of mastery over a piece of the most modern and powerful technology and establish an intimate contact with some of the deepest ideas from science, from mathematics, and from the art of intellectual model building."

Mindstorms. 1980.

Yet sadly Papert's work might be another example of the "Future of Ed-Tech" that hasn't come to pass. He does address this in part in his 1993 book *The Children's Machine*: "Progressive teachers knew very well how to use the computer for their own ends as an instrument of change; School knew very well how to nip this subversion in the bud."

As Bret Victor argues in his keynote: developing new technologies is easy; changing human behaviors, changing institutions, challenging tradition and power is hard.

"Computer-aided inspiration" as Papert envisioned has been mostly trumped by "computer-aided instruction."

Indeed, computer-aided instruction came under development around the same time as Logo and the DynaBook – even earlier, actually. And the history of the future of computer-aided instruction may well tell us more about the ed-tech we've inherited. It certainly points to the ed-tech that many people still want us to have.

The first computer-aided instruction system was PLATO (short for Programmed Logic for Automatic Teaching Operations), a computer system developed at the University of Illinois. 1960 saw the first version, the PLATO I, operate on the university's ILLIAC I computer. Then came PLATO II, PLATO III, and PLATO IV.

The PLATO IV was released in 1972, the same year as Alan Kay's manifesto. It's roughly the same time as Bret Victor situates his "Future of Programming" keynote. Early versions of the PLATO system had a student terminal attached to a mainframe. The software offered mostly "drill and kill" and tutorial lessons. But as the PLATO system developed, new and more sophisticated software was added – more problem-based and inquiry-based lessons, for example. A new programming language called TUTOR enabled "anyone" to create their own PLATO lessons without having to be a programmer. The mainframe now supported multiple, networked computers. Students could communicate with one another, in addition to the instructor. Pretty groundbreaking stuff as this was all pre-Internet.

This networked system made PLATO the site for a number of very important innovations in computing technology, not to mention in ed-tech. Forums, message boards, chat rooms, instant messaging, screen sharing, multiplayer games, and emoticons. PLATO was, as author Brian Dear argues in his forthcoming book *The Friendly Orange Glow* "the dawn of cyberculture."

And again, the familiar story: PLATO's contribution to cyberculture is mostly forgotten.

Arguably PLATO's contribution to ed-tech has been forgotten as well. I'm not sure. I think that we can see in PLATO many of the features in ed-tech today, many of the features that would make Alan Kay and Seymour Papert shudder.

One of the features PLATO boasted: tracking every keystroke that a student made, data on every answer submitted, right or wrong. PLATO offered more efficient computer-based testing. It offered the broadcast of computer-based lessons to multiple locations, where students could work at their own pace. Indeed, by the mid-Seventies, PLATO was serving students in over 150 locations – not just across the University of Illinois campus, but also in elementary schools, high schools, and on military bases. Sensing a huge business opportunity, the Control Data Corporation, the company that built the University of Illinois mainframe, announced that it was going to go to market with PLATO, spinning it out from a university project to a corporate one.

This is where that $500 price tag for Alan Kay's DynaBook is so significant. CDC charged $50 an hour for access to its mainframe, for starters. Each student unit cost about $1900; the mainframe itself cost at least $2.5 million, according to estimates in a 1973 review of computer-assisted instruction. CDC charged $300,000 to develop each piece of courseware.

Needless to say, PLATO as a computer-aided instruction product was a failure. The main success that CDC had with it: selling an online testing system to the National Association of Securities Dealers, a regulatory group that licenses people who sell securities. CDC sold the PLATO trademark in 1989 to The Roach Organization, and it now sells e-learning software under the name Edmentum.

From a machine at "the dawn of cyberculture" to one that delivered standardized testing for stockbrokers. The history of the future of ed-tech. The refrain of this talk: new technologies are easy to develop; new behaviors and new cultures are not.

One final piece of education technology history, this one a little older than the computer-based innovations of the 1960s and 1970s. It's still a machine-based innovation. It's still an object that enables efficient instruction and efficient assessment.

B. F. Skinner's "teaching machine."

I could go back farther than Skinner, admittedly. To a patent in 1866 for a device to teach spelling. Or to a patent in 1909 for a device to teach reading. Or to a patent in 1911 awarded to one Herbert Aikens that promised to teach "arithmetic, reading, spelling, foreign languages, history, geography, literature or any other subject in which questions can be asked in such a way as to demand a definite form of words ... letters ... or symbols."

I could go back to the machine developed by Sidney Pressey. Pressey was psychologist at Ohio State University, and he came up with an idea for a machine to score the intelligence tests that the military was using to determine eligibility to enlistment. Then World War I happened, causing a delay in Pressey's work. He first exhibited his teaching machine at the 1925 meeting of the American Psychological Association. It had four multiple-choice questions and answers in a window, and four keys. If the student thought the second answer was correct, she pressed the second key; if she was right, the next question was turned up. If the second was not the right answer, the initial question remained in the window, and the student persisted until she found the right one. A record of all the student's attempts was kept automatically.

Intelligence testing based on students' responses to multiple-choice questions. Multiple-choice questions with four answers. Sound familiar?

Harvard professor B. F. Skinner claimed he'd never seen Pressey's device when he developed his own teaching machine in the mid 1950s. Indeed, he dismissed Pressey's device, arguing it was a testing and not a teaching machine. Skinner didn't like that Pressey's machine featured multiple-choice questions. His enabled students to enter their own responses by pulling a series of levers. The correct answer made a light go on.

A behaviorist, Skinner believed that teaching machines could provide an ideal mechanism for operant conditioning. "There is no reason why the schoolroom should be any less mechanized than, for example, the kitchen," he argued.

Skinner believed that immediate, positive reinforcement was key to shaping behavior. All human actions could be analyzed this way. Skinner contended that, despite their important role in helping to shape student behavior, "the simple fact is that, as a mere reinforcing mechanism, the teacher is out of date."

Skinner's teaching machine might look terribly out-of-date, but I'd argue that this is the history that still shapes so much of what we see today. Self-paced learning, gamification, an emphasis on real-time or near-real-time corrections. No doubt, ed-tech today draws quite heavily on Skinner's ideas because Skinner (and his fellow education psychologist Edward Thorndike) has been so influential in how we view teaching and learning and how we view schooling.

So much B. F. Skinner. So little Seymour Papert. So little Alan Kay.

I'd argue too that this isn't just about education technology. There's so much Skinner and so little Kay in "mainstream" technology too. Think Zynga, for example. Click, click, click. Level up! Rewards! Achievement unlocked!

As we see our society becoming more and more "technological," it's worth considering the origins and the trajectory of all that tech.

I'll quote Papert here, one more time, to close: "One might say the computer is being used to program the child. In my vision, the child programs the computer, and in doing so, both acquires a sense of mastery over a piece of the most modern and powerful

technology and establishes an intense contact with some of the deepest ideas from science, from mathematics, and from the art of intellectual model building."

May that vision be what guides us forward. May that be what shapes the future of ed-tech.

This keynote was delivered on February 4, 2014 at the EdTechTeacher iPad Summit in San Diego, California. The original version can be found on Hack Education at http://hackeducation.com/2014/02/04/the-history-of-the-future-of-ed-tech/

WORKS CITED

Brian Dear, *The Friendly Orange Glow: The Story of the PLATO System and the Dawn of Cyberculture*. http://friendlyorangeglow.com/

David Greelish, "An Interview with Computing Pioneer Alan Kay." *TIME*. April 2, 2013. http://techland.time.com/2013/04/02/an-interview-with-computing-pioneer-alan-kay/

Alan Kay, "A Personal Computer for Children of All Ages." Proceedings of the ACM Annual Conference. 1971. Vol 1, No. 1.

Seymour Papert, *The Children's Machine: Rethinking School in the Age of the Computer*. Basic Books. 1993.

Seymour Papert, *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books. 1980.

B. F. Skinner, *The Technology of Teaching*. Copley Publishing Group. 1968.

Paula Taylor, *The Kids' Whole Future Catalog*. Random House Books for Young Readers. 1982.

Bret Victor, "The Future of Programming," Dropbox's DBX Conference. July 9, 2013. http://vimeo.com/71278954

2 UN-FATHOMABLE: THE HIDDEN HISTORY OF ED-TECH

A couple of years ago, a friend sent me an exasperated email on the heels of an exclusive technology event he'd attended in Northern California – not in Silicon Valley, but close enough – one with powerful people in the tech industry. Investors. Engineers. Entrepreneurs. Several prominent CEOs of prominent ed-tech startups had been invited to speak there about the state of education – past, present, and future – and their talks, my friend reported, tended to condemn education's utter failure to change or to adopt computing technologies. The personal computing revolution had passed schools by entirely, they argued, and it wasn't until the last decade that schools had started to even consider the existence of the Internet. The first online class, insisted one co-founder of a company that's raised tens of millions of dollars in venture capital since then, was in 2001 at MIT.

And okay, in fairness, these folks are not historians. They're computer scientists, artificial intelligence experts, software engineers. They're entrepreneurs. But their lack of knowledge about the history of education and the history of education technology matters.

It matters because it supports a prevailing narrative about innovation – where innovation comes from (according to this narrative, it comes from private industry, and not from public institutions; from Silicon Valley, that is, not from elsewhere in the world) and when it comes (there's this fiercely myopic fixation on the future).

The lack of knowledge about history matters too because it reflects and even enables a powerful strain in American ideology and in the ideology of the technology industry: that the past is irrelevant, that the past is a monolithic block of brokenness – unchanged and unchanging until it's disrupted by technological innovation, or by the promise of technological innovation, by the future itself.

This ideology shapes the story that many ed-tech entrepreneurs tell about education and about their role in transforming it.

One of my favorite examples of this comes from Salman Khan, the founder of Khan Academy, in a video on "The History of Education" he made with *Forbes* writer Michael Noer in 2012.

It's the history of education "from 1680 to 2050" told in 11 minutes, so needless to say it's a rather abbreviated version of events. It's not titled "The History of Education in the United States," although that would be much better because contributions to education from the rest of the world are entirely absent.

Well, except for the Prussians. Americans involved in education reform and education technology love to talk about the Prussians.

Our current model of education, says Khan, originated at the turn of the nineteenth century: "age-based cohorts" that move through an "assembly line" with "information being delivered at every point."

"This is the Prussian model," the *Forbes* writer Noer adds, "and it's about as inflexible as a Prussian can be." But Khan notes that there were benefits to this as "it was the first time people said, 'No, we want everyone to get an education for free.'"

Then "Horace Mann comes along about 1840" and introduces this concept of free education for everyone to the United States. By 1870, says Khan, public education is

pretty common "but even at that point it wasn't uniform" with different standards and curriculum in different states and cities. So in 1892, "something that tends to get lost in history," a committee of ten – "somewhat Orwellian" quips Noer – meet to determine what twelve years of compulsory public education should look like.

"It was forward looking for 120 years ago," says Noer, "but what's interesting is that we've basically been stuck there for 120 years." Education has been "static to the present day," agrees Khan.

And from 1892, the story they tell jumps ahead, straight to the invention of the Internet – "the mid late Nineties," says Khan as he plots it on his timeline. "The big thing here," says Noer as the two skip over one hundred years or so of history, "is what you've done" with Khan Academy. "One person with one computer can reach millions." This revolutionizes lectures, Noer argues; it revolutionizes homework. "Class time is liberated," adds Khan. This changes everything. Khan Academy (founded in 2006) changes everything that has been stagnant and static since the nineteenth century.

See, this isn't simply a matter of forgetting history – the history of technology or the history of education or the history of ed-tech. It's not simply a matter of ignoring it. It's a rewriting of history, whether you see it as activist or accidental.

To contend, as my friend overheard at that tech event or as Khan implies in his history of education, that schools haven't been involved in the development or deployment of computers or the Internet, for example, is laughably incorrect. It's an inaccurate, incomplete history of computing technology, not simply an inaccurate history of ed-tech.

Take the ILLIAC I, the first von Neumann architecture computer owned by an American university, built in 1952 at the University of Illinois. (The US was beaten by several years by universities here in the UK, I should point out, namely at the nearby University of Manchester.)

Or take PLATO, a computer-based education system, sometimes credited as the first piece of educational computing software, built on the University of Illinois' ILLIAC machine in 1960.

Or take the work of Marc Andreessen, now a powerful figure in Silicon Valley, a venture capitalist with several major investments in ed-tech, who took the work he'd done on the Mosaic Web browser as a student at the University of Illinois in order to start his own company, Mosaic Communications Company, which eventually became the Netscape Communications Company, launching the Netscape Navigator web browser and successfully IPOing in 1995.

The history of education technology is long. The history of education technology is rich. And while it certainly predates Netscape or the von Neumann architecture, the history of education technology is deeply intertwined with the history of computing – and visa versa.

And I could probably stop right there with my keynote. This is really the crux of my message: there's a fascinating and important history of education technology that is largely forgotten, that is largely hidden. It's overlooked for a number of reasons, some of which is wrapped up in the ideologies I've already alluded to.

All this means, if we're going to talk about "Building the Digital Institution: Technological Innovation in Universities and Colleges," the theme of this conference, we probably should know a bit about the history of universities and colleges and

technological innovation and build from there.

Despite all the problems that these institutions have – and good grief, they do have problems – universities and colleges have been the sites of technological innovation. They are the sites of technological innovation. Or they can be. In pockets, to be sure. In spurts, to be sure. Certain developments in certain times in certain places, yes. Certain disciplines making breakthroughs; certain disciplines getting the credit. Certain universities getting the credit for innovating, even when, dare I say, they aren't actually doing anything that new or transformative.

It's not surprising perhaps that the ed-tech startup co-founder in my opening anecdote would credit MIT with offering the first online course. It's one of those universities that consistently gets the credit for "innovation." Perhaps he was thinking of MIT OpenCourseWare which launched in 2002 as an effort to put the university's course materials online in a free and openly licensed format.

A couple of side-notes: 1) That putting course materials online could be confused with offering a course online speaks volumes about this co-founder's startup. 2) This particular ed-tech co-founder attended MIT. 3) Salman Khan is also a MIT graduate, and I think his vision for teaching and learning via a site like Khan Academy draws heavily on that MIT academic culture, where class attendance isn't as important as working through course materials at your own pace with your smartest peers. As long as you can pass the assessments at the end of the course, that's what matters.

It's unlikely, when touting who put classes online first that this ed-tech co-founder from my opening anecdote was thinking of Fathom, the Columbia University-led online learning initiative founded roughly around the date he ascribed to the first "online course." It's unlikely he was thinking of AllLearn, the Stanford, Yale, and Oxford Universities-led online learning initiative of the same period.

Possibly because it's like the movie *Fight Club*. The first rule of the history of online education: you don't talk about Fathom. You don't talk about AllLearn.

And this particular ed-tech startup co-founder certainly wasn't talking about UK e-University, because as with the development of early computers, we ("we Americans," I should qualify here) seem to have forgotten that much has happened outside of the US, let alone outside of Silicon Valley.

Ah, ed-tech of the late 1990s and early 2000s. "The Internet!" as Salman Khan exclaims excitedly.

We don't talk much about that period. We don't talk much about the heady days of the first Dot Com bubble. Have we really forgotten?

It could be that we're reluctant in talking about the first Dot Com bubble because some of us don't want to admit we might just be in the midst of another one. Startups, ed-tech and otherwise, are overhyped and overfunded and overvalued, many with little to show in terms of profit (or educational outcomes).

What's implied by our silence about the Dot Com era perhaps: we know better now than we did then. Or at least the tech is better. Or at least we're not spending as much money to launch startups these days. Or we care more about learning now. Or something.

And yes, some of us simply don't want to talk about the tech and ed-tech failures of the Dot Com era – the failures of Fathom and AllLearn and UKeU and the like – because of the shame of failure. It's not just Silicon Valley entrepreneurs who are at fault

here. I think industry and institutions (particularly elite Ivy League institutions) have buried those failures. That's a pity since there's much to learn.

I realize that most of the folks here know these stories, but I'm going to repeat them anyway.

Fathom opened in 2000 and closed in 2003.

AllLearn opened in 2001 and closed in 2006.

UKeU opened in 2003 and closed in 2004.

$30 million invested into the Fathom initiative by Columbia University.

$12 million was invested into AllLearn from various schools and foundations.

£62 million was earmarked for and £50 million was spent by the British government on UKeU.

For a little comparison: edX launched in 2012 with an initial $60 million investment from Harvard and (yes) MIT. Coursera launched in 2012 with a total venture capital investment of $85 million. Udacity launched in 2012 with a total (disclosed) venture capital investment of $20 million.

This notion that it's easier and cheaper to launch a startup in the 2010s, that thanks to open source technologies and the cloud and the like that we needn't funnel so much money into ed-tech startups. Well...

Thanks to the Internet Archive Wayback Machine, we can see what Fathom and AllLearn's websites looked like circa 2001. It's an important tool as if you search for "UKeU" today, you might accidentally stumble upon Ukulele University. If you compare these sites to contemporary online education sites like Coursera or FutureLearn, you can see some changes – improvements no doubt – in Web design. But what's really changed in the decade or so between the Dot Com-era online courses and today's versions?

What's changed in terms of institutional involvement? What's changed in terms of branding? What's changed in terms of course content, and what's changed in terms of the "ed-tech" under the hood? What hasn't changed? What's the same?

The course content for Fathom and AllLearn was similar to what we see being offered online today. That's not a surprise, as such is the makeup of the typical college course catalog: a broad swath of classes in science, technology, humanities, professional development, business, and law. Some 2000 courses were offered via Fathom. There were 110 offered on AllLearn. 25 on UK e-University. (Is that correct?!) There are over 500 courses offered via Coursera.

The technology hasn't changed much in the intervening decade. (And the phrase "content delivery system" is still used to describe online education, sadly.) The Dot Com era courses offered "primary source documents, animations, interactive graphics, audio slide shows, and streaming videos." Today's online courses look much the same, and despite their boasts about better assessment tools – automated essay graders and the like – multiple choice quizzes, a historical artifact from the earliest teaching machines of the 20th century, still dominate.

The marketing pitch to students hasn't changed much either: "Online courses from the world's best universities" – that's the tagline on the edX site. The "world's best courses" – that's what Coursera promises. "Enjoy free online courses from leading UK and international universities" – that's FutureLearn's promise. The "world's most trusted sources of knowledge" –that was Fathom's. The focus, then and now, is on the prestige of the institutions involved. And they are some of the very same institutions. Stanford. Yale.

Columbia.

AllLearn, short for the Alliance for Lifelong Learning, stressed that its classes were just that: an opportunity for continuing education and lifelong learning. Udacity stresses something different today: it's about "advancing your career." It's about "dream jobs."

There's been plenty of hype about these new online platforms displacing or replacing face-to-face education, and part of that does connect to another powerful (political) narrative: that universities do not adequately equip students with "21st century skills" that employers will increasingly demand. But by most accounts, those who sign up for these courses still fall into the "lifelong learner" category. The majority has a college degree already.

The question remains unresolved, a decade later, as to whether or not people will pay for these online courses (or for certification after successful completion) to such an extent that these online initiatives can ever become financially sustainable, let alone profitable. That's even accounting for the massive increase since the early 2000s in the cost of higher education (in the US and now elsewhere) alongside the growing demand for everyone to have some sort of college credential.

From a 2002 *New York Times* article about universities' efforts to move online, "Lessons Learned at Dot Com U": "college campuses and dot-coms had looked at the numbers and anticipated a rising tide of enrollment based on baby boomers and their children as both traditional students and those seeking continuing education. In short, the colleges essentially assumed that if they built it, students would come."

"We hope it's enough money to get us to profitability,'" Coursera co-founder Daphne Koller told *The New York Times* in the summer of 2013 when her company announced it had raised another $43 million. "We haven't really focused yet on when that might be." Echoing the *Field of Dreams* reference from a decade earlier – that's a baseball movie reference, a terrible thing to invoke in a keynote in the UK, I realize: if you build it, they will come. Indeed, Koller has admitted that her investors have told her that if you do the "right thing" in education, the profits will follow.

Perhaps they will.

We can see already the pressures for Coursera to find a path to profitability. It has raised $85 million in venture capital after all, not in university endowment or in foundation funding. In recent months, Coursera has shuffled its executive team quite a bit, adding a venture capitalist from fabled investment firm Kleiner Perkins Caufield and Byers as President and adding a former Yale President as CEO. Co-founder Andrew Ng has stepped away from day-to-day work at the company, although he remains Chairman of the Board.

The new CEO of Coursera, Richard Levin, as it just so happens, was at the helm at Yale in the AllLearn era. (He was the chair of AllLearn as well.) One might assume then that he must have a significant amount of expertise and much wisdom gleaned from the university's Dot Com era ed-tech ventures. Levin, an economist by training, must know a bit about the history of education and the history of technology and the history of ed-tech. Or at least he should know a bit about the history of the economics of ed-tech. Right?

In an interview with *The New York Times* this spring, Levin offered this explanation as to why AllLearn did not succeed: "It was too early. Bandwidth wasn't

adequate to support the video. But we gained a lot of experience of how to create courses, and then we used it starting in 2007 to create very high quality videos, now supported by adequate bandwidth in many parts of the world, with the Open Yale courses. We've released over 40 of them, and they gained a wide audience."

AllLearn failed, he argues, because of bandwidth. Bandwidth.

"The Internet bandwidth in most homes was inadequate for properly sharing course material," Levin contends. Actually, AllLearn offered its materials via CD-ROM as well, and like many sites in that period, AllLearn recognized that streaming video content might be challenging for many users. It allowed them to turn off some of the high-bandwidth features and download rather than watch video online.

Remember too, AllLearn was marketed as a "lifelong learning" site. Its pitch was to alumni of the elite universities involved as well as to the general public. The former would pay about $200 per course; the latter about $250. (One creative writing class charged $800 in tuition.) So are we to believe that those groups – alumni and keen lifelong learners – were unable to access AllLearn due to bandwidth issues? That they'd balk at paying for good Internet but not balk at the AllLearn fees? This is an assertion, an explanation that my colleague Mike Caulfield has questioned: "All-Learn folded in 2006, when broadband was at a meager 20% adoption. Today, it's different, supposedly. It's at 28%. Are we to really believe that somewhere in that 8% of the population is the difference between success and failure?" asks Caulfield.

Caulfield also questions what Levin learned from OpenYale, the ed-tech venture that followed the demise of AllLearn. By Caulfield's calculations, those courses were created using "$4 million dollars of Hewlett money. And the videos are basically recordings of class lectures. Four million dollars for forty filmed courses, or, if you prefer, $100,000 a course for video lectures."

That's close to the cost for course production you hear bandied about today by professors who've created Coursera classes, for what it's worth.

It's this discrepancy between the costs and the revenue, an inability to find a sustainable business model that plagued the Dot Com era online initiatives. From a 2003 article in the Columbia student newspaper: "Fathom spent money at an unsustainable rate. In 2001, Fathom burned through almost $15 million, and generated revenues of only $700,000." And this is what plagues Coursera today.

This is (in part) why history matters. Well, history and a bit of humility, I'd add. It's not easy to reflect on our failures – the failures of Dot Com era ed-tech in this case – and move forward; but that's how we make progress.

It's important too to recognize the successes of the Dot Com era and to remember that, despite the failures of initiatives like AllLearn and Fathom, there were many online education programs founded in roughly the same period that didn't fold and that went on to be sustainable. Many of these continue to operate today.

I'd argue however that (sadly) that one of the most significant successes of the Dot Com ere – financial successes, that is – is one that has left an indelible mark on ed-tech. And that's the success of the learning management system. The technology, the industry.

While learning management system software predates the Internet, it was the Internet that became its big selling point. From *The Washington Post* in 1999: "Blackboard Chalks Up a Breakthrough; Its Educational Software Lets Colleges Put

Classes on the Internet." (Several years, I'd like to point out, prior to the date in my opening anecdote when MIT supposedly offered the first course online.)

The LMS – or the VLE, I should say while here in the UK – has profoundly shaped how schools interact with the Internet. The LMS is a piece of administrative software. There's that word "management" in there that sort of gives it away for us in the US at least: that this software that purports to address questions about teaching and learning but that really works to "manage" and administer, in turn often circumscribing pedagogical possibilities. You can see its Dot Com roots too in the LMS functionality and in its interface. I mean, some LMSes still look like software from the year 2000! The LMS acts as an Internet portal to the student information system, and much like the old portals of the Dot Com era – much like AOL for example – it cautions you when you try to venture outside of it. You can access the LMS through your web browser but it is not really of the web.

The learning management system is a silo, a technological silo, by design. This isn't because the technology isn't available to do otherwise. Rather, it's a reflection of the institution of education. The LMS silo works because we tend to view each classroom as a closed entity, because we view each subject or discipline as atomistic and distinct. Closed. Centralized. Control in the hands of administrators, teachers, and IT but rarely in the hands of learners.

If you look at the much-hyped online courses of today – those offered on the Coursera or the edX platforms, for example – you can see the influence of the LMS. Each course you enroll in is separate, siloed. At the end of the term, your access to your course disappears. There's a tab on the LMS so you can navigate to the syllabus and a tab for assignments and one for assessments, and there is, of course – thanks early Internet technology! – a discussion forum. A message board. It isn't an accident, and it certainly isn't an innovation, that our online classes look this way.

It doesn't have to look this way, of course. There are other stories we could tell about education technology's past; there are other paths forward. Again, there's this hidden history of ed-tech (and of computer tech as well), and it's worth considering why so much has been forgotten or overlooked or dismissed. Ted Nelson. Douglas Englebart. Or the person I always point to: Seymour Papert.

Computers, argued Papert, should unlock children's "powerful ideas." That's the subtitle to his 1980 book *Mindstorms*, a book that I insist people in ed-tech read (although admittedly Papert's work is geared towards younger children rather than adult learners). *Mindstorms* addresses: "how computers can be carriers of powerful ideas and of the seeds of cultural change, how they can help people form new relationships with knowledge that cut across the traditional lines separating humanities from sciences and knowledge of the self from both of these. It is about using computers to challenge current beliefs about who can understand what and at what age. It is about using computers to question standard assumptions in developmental psychology and in the psychology of aptitudes and attitudes. It is about whether personal computers and the cultures in which they are used will continue to be the creatures of 'engineers' alone or whether we can construct intellectual environments in which people who today think of themselves as 'humanists' will feel part of, not alienated from, the process of constructing computational cultures."

Computers, Papert insisted, will help children gain "a sense of mastery over a piece of the most modern and powerful technology and establish an intimate contact with

some of the deepest ideas from science, from mathematics, and from the art of intellectual model building."

But as we see with the LMS, ed-tech has come to mean something else. As Papert notes in his 1993 book *The Children's Machine*: "Progressive teachers knew very well how to use the computer for their own ends as an instrument of change; School knew very well how to nip this subversion in the bud."

"Computer-aided inspiration," as Papert envisioned, has been mostly trumped by "computer-aided instruction."

And we come full circle now to a technology I mentioned in passing at the beginning of my talk: PLATO, Programmed Logic for Automatic Teaching Operations, a computer system developed at the University of Illinois in the 1960s on its ILLIAC machine.

Early versions of the PLATO system had a student terminal attached to a mainframe. The software offered mostly "drill and kill" and tutorial lessons. But as the PLATO system developed, new and more sophisticated software was added. There were more problem-based lessons, for example. A new programming language called TUTOR enabled "anyone" to create their own PLATO lessons without having to be a programmer. The mainframe came to support multiple, networked computers. Students could communicate with one another, in addition to the instructor. And this was all pre-Internet, pre-Web.

This networked system made PLATO a site for a number of very important innovations in computing technology, not to mention in ed-tech. Forums, message boards, chat rooms, instant messaging, screen sharing, multiplayer games, and emoticons. PLATO was, as author Brian Dear argues in his forthcoming book *The Friendly Orange Glow* "the dawn of cyberculture."

But as with so much ed-tech history, PLATO's contribution to cyberculture is mostly forgotten. Yet we can still see remnants of PLATO in many of the features in ed-tech today, including of course, the learning management system. And if the learning management system has trapped us in a moment of Dot Com era tech – trapped in the old Internet portal –it may be that ed-tech's roots in PLATO have trapped us in an old "mainframe" mindset as well.

See, there are numerous legacies here. One of the features PLATO boasted: tracking every keystroke that a student made, data on every answer submitted, right or wrong. Sound familiar? PLATO offered more efficient computer-based testing. Sound familiar? It offered the broadcast of computer-based lessons to multiple locations, where students could work at their own pace. Sound familiar? Indeed, by the mid-Seventies, PLATO was serving students in over 150 locations – not just across the University of Illinois campus, but also in elementary schools, high schools, and on military bases. Sensing a huge business opportunity – the notion of tapping into the giant "education market" is not new – the Control Data Corporation, the company that built the University of Illinois mainframe, announced that it was going to go to market with PLATO, spinning it out from a university project to a corporate one.

CDC charged $50 an hour for access to its mainframe, for starters. Each student unit cost about $1900; the mainframe itself $2.5 million, according to some estimates. CDC charged $300,000 to develop each piece of courseware. (So okay, I guess it is getting a little cheaper to develop courseware.)

Needless to say, PLATO as a commercialized computer-aided instruction product was largely a failure. The main success that CDC had with it: selling an online testing system to the National Association of Securities Dealers, a regulatory group that licenses stockbrokers.

Yet like the learning management system, the idea of computer-assisted instruction has retained an incredibly powerful hold over ed-tech. Indeed, as the history of PLATO shows us, the two are interconnected. Computer-based instruction. Computer-based management.

As we move forward, "building the digital institution," I think we must retrace and unwind some of these connections.

Why are we building learning management systems? Why are we building computer-assisted instructional tech? Current computing technologies demand neither. Open practices don't either. Rather, it's a certain institutional culture and a certain set of business interests that do.

What alternatives can we build? What can we imagine? Can we envision a future of learner agency, of human capacity, of equity, of civic responsibility, of openness for example?

I called this talk "Un-Fathom-able," thumbing my nose I confess at the failures of Fathom and what I think we may soon see as the failure of Coursera. I called this talk "Un-Fathom-able" too because I fear that there's much in ed-tech that we've failed to explore – partly, I would argue, that's because we have failed to learn and to reflect on the history of ed-tech. It's easy to blame technologists, I suppose. But I think all this runs deeper than that. There's been a failure of imagination to do something bold and different, something that, to borrow Papert's phrasing, unlocks "powerful ideas" in learners rather than simply re-inscribing powerful institutional mandates.

We can't move forward until we reconcile where we've been before.

This keynote was delivered on June 18, 2014 at CETIS in Bolton, UK. The original transcript, along with the slides for this talk, can be found on Hack Education at http://hackeducation.com/2014/06/18/unfathomable-cetis2014/

WORKS CITED

Yoni Appelbaum, "Digitalia Columbiana." *Columbia Daily Spectator*. January 30, 2003. http://columbiaspectator.com/2003/01/30/digitalia-columbiana

Mike Caulfield, "Experience Without Humility Is Not Very Helpful At All." April 16, 2014. http://hapgood.us/2014/04/16/experience-without-humility-is-just-experience/

Brian Dear, *The Friendly Orange Glow: The Story of the PLATO System and the Dawn of Cyberculture*. http://friendlyorangeglow.com/

D. D. Guttenplan, "Out in Front, and Optimistic, About Online Education." *The New York Times*. April 13, 2014. http://www.nytimes.com/2014/04/14/education/out-in-front-

and-optimistic-about-online-education.html

Katie Hafner, “Lessons Learned at Dot-Com U.” *The New York Times*. May 2, 2002. http://www.nytimes.com/2002/05/02/technology/lessons-learned-at-dot-com-u.html

Salman Khan and Michael Noer, “The History of Education.” http://youtu.be/LqTwDDTjb6g

Tamar Lewin, “Coursera, an Online Education Company, Raises Another $43 Million.” *The New York Times*. July 10, 2013. http://bits.blogs.nytimes.com/2013/07/10/coursera-an-online-education-company-raises-another-43-million/

Seymour Papert, *The Children’s Machine: Rethinking School in the Age of the Computer*. Basic Books, 1993.

Seymour Papert, *Mindstorms: Children, Computers, and Powerful Ideas*. Basic Books. 1980.

3 TEACHING MACHINES: A BRIEF HISTORY OF "TEACHING AT SCALE"

Having spent the last week outside of the United States, in Europe talking about education technology, I have been reminded how much context matters. It matters when we talk about education technology's future, its present, and its history. Despite all the talk about the global economy, global communications, global democratization of education, context matters. The business and the politics and the stories of ed-tech are not universal.

I think that's something for you to keep in mind as you work your way through this course. It's something to think about when we start to imagine and to build "education at scale." What happens to context? What happens to local, regional education – its history, its content (the curriculum if you will), its cultural relevance and significance, and finally its politics, its practices?

How does technology shape this? How might technology erase or ignore context?

What ideologies does education technology carry with it? Do these extend, reinforce, or subvert existing ideologies embedded in education?

Because of the forward-facing ideology of technology – that is, its association with progress, transformation, "the future" – I think we do tend to forget its history. We tend to ignore its ideology. I think that dovetails quite powerfully too with parts of American ideology and identity: an emphasis on and excitement for "the new"; a belief that this country marked a formal break from other countries, from other histories. A belief in science and business and "progress."

Yesterday was one of those regularly scheduled moments when the technology industry puts all that ideology on display: an Apple keynote, where new products are introduced that have everyone cooing about innovation, that have everyone prepared to declare last year's hardware and software obsolete, and often that have education technology writers predicting that new Apple products are going to revolutionize the way we teach and learn.

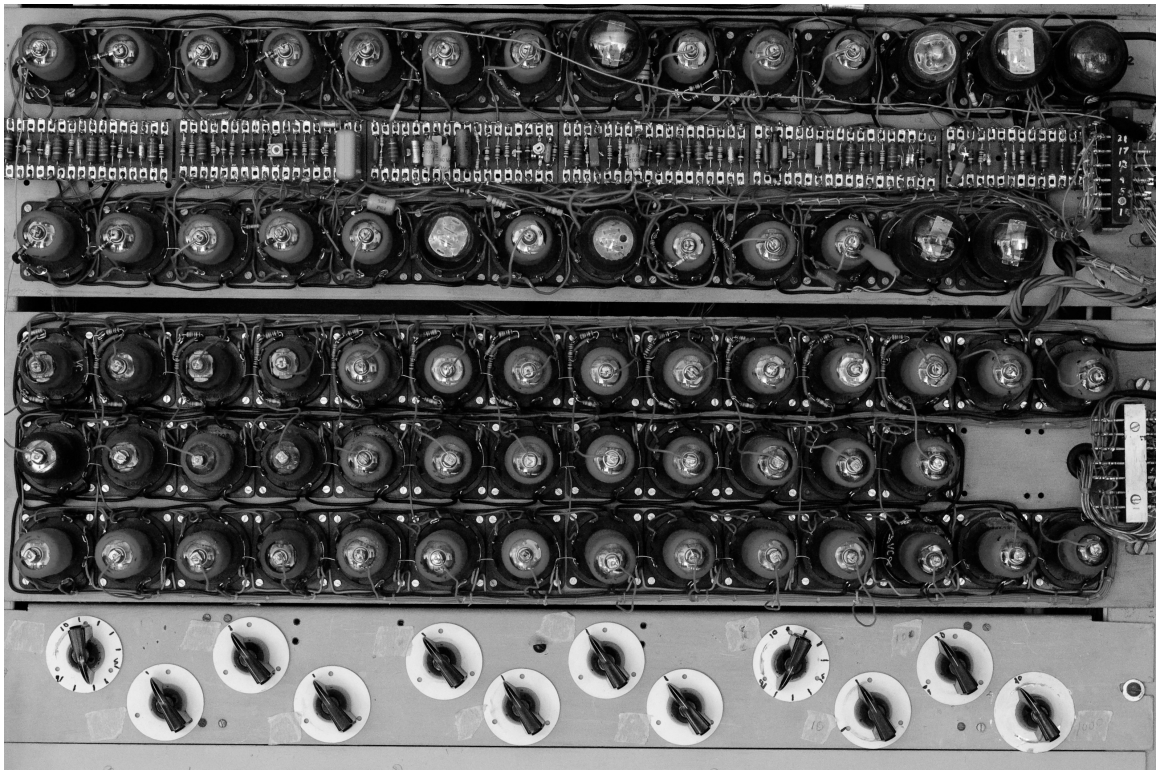

This image is not from the guts of the Apple Watch, of course, or the new iPhone. It is a close-up of the (rebuild of the) Colossus, the world's first electronic, programmable computer. The Colossus prototype was built at Bletchley Park, site of the British government's Code and Cypher School during World War II. It was used to help successfully decrypt German military communications.

Like I said, ideology is embedded in technology. Computers' origins are wrapped up in war and cryptography and surveillance. How does that carry forward into education technology?

When we talk about "education technology" we do tend to focus on the things that teachers and students can do with computers. But education technology certain pre-dates the Colossus (1943). And perhaps we could reach as far back to Plato's *Phaedrus* to see the sorts of debates about what the introduction of new technologies – in this classic example, Socrates' skepticism about the technology of writing – would do to education and more broadly, to culture.

I'm in the middle of writing a book called *Teaching Machines*, a cultural history of the science and politics of ed-tech. An anthropology of ed-tech even, a book that looks at knowledge and power and practices, learning and politics and pedagogy. My book explores the push for efficiency and automation in education: intelligent tutoring systems, artificially intelligent textbooks, robo-graders, and robo-readers.

This involves, of course, a nod to "the father of computer science" Alan Turing, who worked at Bletchley Park, and his profoundly significant question "Can a machine think?"

I want to ask in turn, "Can a machine teach?"

Then too: Why would we want a machine to teach? What happens, as this course is asking you to consider, when we use machines to teach and learn "at scale"?

And to Turing's question, what will happen to humans when (if) machines do "think"? What will happen to humans when (if) machines "teach"? What will happen to labor and what happens to learning?

And, what exactly do we mean by those verbs "think" and "teach"? When we see signs of thinking or teaching in machines, what does that really signal? Is it that our machines are becoming more "intelligent," more human? Or is it that humans are becoming more mechanical?

There's a tension there between freedom and standardization and mechanization that both technology and education grapple with.

Rather than speculate about the future, I want to talk a bit about the past.

I want to suggest that the history of education in the US (and again, this is why context really matters) is woven incredibly tightly with the development of education technologies, and specifically the development of teaching machines. Since the mid-nineteenth century, there have been a succession of technologies that were supposed to improve, if not entirely transform, the way in which teaching happened: the chalkboard, the textbook, radio, film, television, computers, the Internet, the Apple Watch, and so on.

There are a number of factors at play here that make education so susceptible to the technological influence. US Geography, for example: how do you educate across great distances? National identity: what role should schools play in enculturation, in developing a sense of American-ness? Should curriculum be standardized across the country? If so, how will that curriculum be spread? Individualism: how do we balance the desire to standardize education with our very American belief in individualism? How do we balance "mass education" with "meritocracy"? How do we rank and rate students? Industrialization: what is the relationship between schools and business? Should businesses help dictate what students should learn? Should schools be run like businesses? Can we make school more efficient?

These questions – how we've asked and answered them – shape the ways in which education technology has been developed and wielded. Despite what you often hear that technologies will change teaching and learning, more likely technologies re-inscribe the traditional functions and practices of education.

For those interested in the history of education technology, I recommend Larry Cuban's book *Teachers and Machines*. He's better known for his book *Oversold and Underused*, which looks at computers and schools, but *Teachers and Machines* is interesting because you see the tension around technology in general. It's not simply teachers' reluctance to adopt computers, that is. His book looks at attempts to bring film (in the 1910s), radio (in the 1920s), and television (in the 1960s) into the classroom. These are all broadcast technologies, obviously. They're designed to "deliver educational content," a phrase I really hate. And they were all met with a certain amount of resistance.

Despite an emphasis on "content delivery systems," that's not to say that there weren't some really fascinating projects and predictions in the twentieth century. Take, for example, the Midwest Program on Airborne Television Instruction which operated two DC-6 aircraft out of Purdue University Airport using a technology called "Stratovision" to broadcast educational television to schools, particularly to those who couldn't otherwise pick up a TV signal.

I think there are some key lessons to be learned from these broadcast technologies. I think they're lessons that the MOOC providers, whose marketing sounds an awfully lot like some of these twentieth century "innovators," could do well to learn from. If nothing else, how much are we still conceptualizing technologies that "deliver content" and "expand access"? How does "broadcast" shape what we mean when we talk about "scaling" our efforts? How does "broadcast" fit neatly into very old educational practices centered on the teacher and centered on the content?

You could argue that film and radio and airborne television are "teaching machines," but typically the definition of "teaching machines" involves more than just "content delivery." It involves having an instructional and an assessment component as well. But again, these devices have a very long history that certainly predates computers. The earliest known patent in the United States was issued in 1809 to H. Chard for a "Mode of Teaching to Read." The following year S. Randall filed a patent entitled "Mode of Teaching to Write." Halcyon Skinner (no relation to Harvard psychology professor B. F. Skinner) was awarded a patent in 1866 for an "Apparatus for Teaching Spelling." The machine contained a crank, which a student would turn until he'd arranged the letters to spell the word in the picture. The machine did not, however, give the student any feedback if it was right or wrong.

Between Halcyon Skinner's 1866 teaching machine and the 1930s, there were an estimated 600 to 700 patents filed on the subject of teaching and schooling. The vast majority of these were filed by inventors outside of the field of education. Halcyon Skinner, for example, also filed for a patent for a "motor truck for cars," "tufted fabric," a "needle loom," a "tubular boiler," and many other inventions.

There's some debate about whether or not these early devices "count" as teaching machines, as they don't actually do all the things that education psychologists later decided were key: continuous testing of what students are supposed to be learning; immediate feedback on whether a student has an answer correct; the ability for students to "move at their own pace"; automation.

American psychologist Sidney Pressey is generally credited with being the first person whose machine met all these requirements. He displayed a "machine for intelligence testing" at the 1924 meeting of the American Psychological Association. Pressey received a patent for the device in 1928.

His machine contained a large drum that rotated paper, exposing a multiple-choice question. There were four keys, and the student would press the number that corresponded to the right answer. Pressey's machine had two modes of operation: one labeled "test" and the other labeled "teach." In test mode, the machine would simply record the responses and calculate how many were correct. In teaching mode, the machine wouldn't proceed to the next question until the student got the answer right. The machine did still track how many keys were pressed until the student got it correct. You could also add an attachment to the machine that was essentially a candy dispenser. It allowed the experimenter to set what Pressey called a "reward dial," determining the number of correct responses required to receive a candy reward. Once the response criterion had been reached, the device automatically delivered a piece of candy to a container in front of the subject.

For a prototype that was converted from a sewing machine, we can see in Pressey's machine so much about 20th century education theory and practice – and so

much of that that's still with us today. There's the connection to intelligence testing and the First World War and a desire to create a machine to make that process more standardized and efficient. There's the nod to the work of education psychologist Edward Thorndike: his laws of recency and frequency that dictated how students were supposed to move through material. There is the four answer multiple-choice question. How much of this is now "hard coded" into our education practices? How much of this is now "hard coded" into our education technology?

Sidney Pressey tried very hard to commercialize his teaching machines, but without much success. It wasn't until a few decades later that the idea really took off. And as such, "teaching machines" are probably most closely associated with the work of B. F. Skinner. (He did not receive the patent for his device until 1961.)

Skinner came up with the idea for his teaching machine in 1953. Visiting his daughter's fourth grade classroom, he was struck by its inefficiencies. Not only were all the students expected to move through their lessons at the same pace, but when it came to assignments and quizzes, they did not receive feedback until the teacher had graded the materials – sometimes a delay of days. Skinner believed that both of these flaws in school could be addressed through a machine, and built a prototype that he demonstrated at a conference the following year.

All these elements were part of Skinner's teaching machines: the elimination of inefficiencies of the teacher, the delivery of immediate feedback, the ability for students to move through standardized content at their own pace.

Today's ed-tech proponents call this "personalization."

Teaching – with or without machines – was viewed by Skinner as reliant on a "contingency of reinforcement." The problems with human teachers' reinforcement, he argued, were severalfold. First, the reinforcement did not occur immediately; that is, as Skinner observed in his daughter's classroom, there was a delay between students completing assignments and quizzes and their work being corrected and returned. Second, much of the focus on behavior in the classroom has to do with punishing students for "bad behavior" rather than rewarding them for good.

As Skinner wrote in his book *Beyond Freedom and Dignity*, "We need to make vast changes in human behavior. . . . What we need is a technology of behavior." Teaching machines are one such technology.

Skinner's teaching machine differed from Pressey's in that it did not have students push on buttons to respond to multiple-choice questions. Students had to manually enter their own answers. Skinner felt it was important that students could formulate their own responses. And he worried too that selecting the wrong answer on a multiple-choice question was the wrong sort of behavioral reinforcement.

As with Pressey's teaching machines, we can see in Skinner's some of these elements that still exist in our technologies today. Behaviorism in general: excitement about gamification and "nudges" and notifications from our apps all designed to get us to "do the right thing" (whatever that means). And we see too a real excitement about the potential for transforming classrooms with gadgetry.

"There is no reason," Skinner insisted, "why the schoolroom should be any less mechanized than, for example, the kitchen." Indeed in the 1960s, there was a huge boom in teaching machines. There were door-to-door teaching machine salesmen, I kid you not, including those who sold the Min-Max made by Grolier, the encyclopedia company.

But alongside the excitement were the fears about robots teaching the children. And the machines were expensive, as was the development of the "programmed instruction" modules.

So excitement faded, just as new devices started to be developed – ones that were computer-based. Ones that promises "intelligence."

Intelligence, along with all the promises that teaching machines have made for a century now: efficiency, automation, moving at your own pace, immediate feedback, personalization.

Thomas Edison predicted in 1913 that textbooks would soon be obsolete. In 1962, *Popular Science* predicted that by 1965, over half of students would be taught by machines. I could easily find similar predictions made today about MOOCs or adaptive technology or Apple Watches. These themes persist, and it's worth asking why.

I think you can explain a lot of it when you look at history and think about ideology – what we bring into our technologies, what we ask them to do, and how and why.

This talk was given to the Harvard Graduate School of Education class "The Future of Learning at Scale" on September 10, 2014. The original transcript can be found on Hack Education at http://hackeducation.com/2014/09/10/teaching-machines-teaching-at-scale/

WORKS CITED

Larry Cuban, *Oversold and Underused: Computers in the Classroom*. Harvard University Press, 2003.

Larry Cuban, *Teachers and Machines: The Classroom Use of Technology Since 1920*. Teachers College Press, 1986.

B. F. Skinner, *Beyond Freedom and Dignity*. Hackett Publishing Company, 1971.

B. F. Skinner, *The Technology of Teaching*. Copley Publishing Group. 1968.

"Teaching Machines: Do They Or Don't They?" *Popular Science*. 1962. Vol. 181, No. 6.

Alan Turing, "Computing machinery and intelligence." *Mind*. 1950. Vol. 59, 433-460.